

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**AN INTERFACE AND RELATED METHODS FACILITATING
MOTION COMPENSATION IN MEDIA PROCESSING**

Inventor(s):
Gary Sullivan
Robin Speed
William C. Powell
Nicholas P. Wilt
Chad Fogg

ATTORNEY'S DOCKET NO. MS1-600US

100243-08362600

1 **RELATED APPLICATIONS**

2 This application claims priority to a provisional application entitled *An*
3 *Adaptive Multimedia Application Interface*, serial number 60/198,938, filed on
4 April 21, 2000 by Sullivan, et al. and commonly assigned to the assignee of the
5 present invention.
6

7 **TECHNICAL FIELD**

8 This invention generally relates to video processing and, in particular, to a
9 multimedia application program interface (API) that automatically identifies and
10 dynamically adapts to processing system capability to improve multimedia
11 processing performance.
12

13 **BACKGROUND OF THE INVENTION**

14 With recent improvements in processing and storage technologies, many
15 personal computing systems now have the capacity to receive, process and render
16 multimedia objects (e.g., audio, graphical and video content). The multimedia
17 content may be delivered to the computing system in any of a number of ways
18 including, for example, on a compact disk read-only memory (CD-ROM), a digital
19 versatile disk read-only memory (DVD-ROM), via a communicatively coupled
20 data network (e.g., Internet), and the like. Due to the amount of data required to
21 accurately represent such multimedia content, it is typically delivered to the
22 computing system in an encoded, compressed form. To render the multimedia, it
23 must be decompressed and decoded before it is communicated to a display and/or
24 audio device.
25

A number of multimedia standards have been developed that define the format and meaning of encoded multimedia content for purposes of distribution. Organizations such as the Moving Picture Experts Group (MPEG) under the auspices of the International Standards Organization (ISO), and the Video Coding Experts Group (VCEG) under the auspices of the International Telecommunications Union (ITU), have developed a number of multimedia coding standards, e.g., MPEG-1, MPEG-2, MPEG-4, H.261, H.263, and the like. Such standards define the format and meaning of the coded multimedia content, but not how the encoded content is to generated, and only defines the decoding process in mathematical terms. Consequently, a number of hardware and software solutions have been developed by a number of companies to encode, decode and render multimedia content, often employing proprietary techniques to recover the multimedia content from a particular standardized format.

Simplistically speaking, the encoding process removes spatial and temporal redundancies from the media content, thereby reducing the amount of data needed to represent the media content and, as a result, reducing the bandwidth burden to store and/or transmit such media content. A common encoding process includes a digitization/filtering stage, a prediction stage, and a transformation and difference coding stage. In the digitization/filtering stage, the received analog media content is digitized using, for example, an analog to digital converter and is filtered to remove artifacts. In the prediction stage, spatial and temporal redundancies are identified and removed/reduced using motion estimation prediction techniques. The transformation and difference coding process involves a transformation filtering step (e.g., Discrete Cosine Transform (DCT)), followed by a quantization step and entropy encoding.

1 interface with particular decoder applications. In this regard, prior art solutions
2 often required the accelerator hardware vendors to develop an API to interface
3 their board with any of a plurality of decoder applications that an end-user may
4 employ to control and render multimedia content.

5 As introduced above, however, each manufacturer of multimedia decoding
6 applications/accelerators has taken an individual proprietary approach to decoding
7 multimedia content. That is, each of the decoder applications and multimedia
8 accelerators available in the market offer different levels of functionality, often
9 utilizing different data formats or APIs to expose the same basic capability. One
10 accelerator may provide the inverse transformation (e.g., IDCT) as well as motion
11 compensated prediction capability, while another (perhaps lower-end) multimedia
12 accelerator will rely on the host-based decoder application to perform the inverse
13 transformation process and merely provide the motion compensated prediction
14 and/or display format conversion. Consequently, each decoder
15 application/multimedia accelerator combination is a unique multimedia processing
16 system, which heretofore has required a dedicated API.

17 Another negative consequence of the API proliferation associated with each
18 multimedia accelerator is that it is often necessary or desirable to make changes to
19 the multimedia accelerator – improve processing capability, alter processing
20 techniques, accommodate processing improvements, accommodate developments
21 in computing system technology, etc. Heretofore, whenever such changes were
22 made to the accelerator, a change was necessitated in one or more of the API's
23 associated with the accelerator. In addition to the increased likelihood for the
24 proliferation of unnecessary API's in the end-user's computing system (which
25 may adversely affect system performance), this also unnecessarily complicates the

task of writing a decoder application which is intended to use the acceleration capabilities, potentially rendering the decoder incompatible with some accelerators.

Thus, an adaptive multimedia application program interface that transcends particular software and hardware characteristics is needed, unencumbered by the above limitations commonly associated with the prior art.

SUMMARY OF THE INVENTION

This invention concerns a multimedia application program interface (API) facilitating the use of any one or more of a plurality of multimedia accelerators with a decoder application. According to a first implementation of the present invention, a method facilitating media processing is presented comprising generating a motion compensated prediction of a region of encoded media content, receiving an indication of a quantity of remaining residual samples for refining the prediction, on a per-region basis, and adding a quantity of such samples to the prediction to generate a modified prediction, and subtracting a quantity of such samples from the modified prediction, when so indicated.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of an example computer system incorporating the teachings of the present invention;

Fig. 2 is a block diagram of an example multimedia application program interface (API) incorporating the teachings of the present invention, according to one implementation of the present invention;

1 Figs. 3 and 4 provide a graphical illustration of an example control
2 command data structure and a residual difference data structure, respectively,
3 according to one aspect of the present invention;

4 Fig. 5 is a flow chart of an example method interfacing any decoder
5 application with any accelerator without *a priori* knowledge of the decoder or
6 accelerator to be used, according to one implementation of the present invention;

7 Fig. 6 is a flow chart of an example method of decoding media content,
8 according to one example implementation of the present invention;

9 Fig. 7 is a flow chart of an example method facilitating host-based entropy
10 decoding, according to one aspect of the present invention;

11 Fig. 8 is a flow chart of an example method facilitating application control
12 of an accelerator deblocking filter, in accordance with one aspect of the present
13 invention;

14 Fig. 9 is a block diagram of an example multimedia API, according to an
15 alternate implementation of the present invention; and

16 Fig. 10 is a block diagram of an example storage medium comprising a
17 plurality of executable instructions that when executed implement the multimedia
18 API of the present invention, according to one embodiment of the present
19 invention.

20 21 **DETAILED DESCRIPTION**

22 This invention concerns an application program interface (API) that
23 dynamically adapts to the processing capability of a multimedia processing system
24 to improve multimedia processing performance. In this regard, the present
25 invention is an enabling technology that facilitates innovation in multimedia

processing (e.g., encoding and decoding of media content). For ease of illustration and explanation, and not limitation, the teachings of the present invention will be developed within the implementation context of a video decoding system. As such, certain aspects of video decoding process(es) will be described in the context of the present invention. Thus, it is expected that the reader be generally familiar with multimedia decoding. In particular, familiarity with one or more of the H.261, MPEG-1, H.262/MPEG-2, H.263, and MPEG-4 standards will be useful in understanding the operational context of the present invention:

ITU-T Recommendation H.261: Video Codec for Audiovisual Services at
Px64 kbit/s, 1993.

ISO/IEC 11172-2 (MPEG-1 Video): Information technology -- Coding of moving pictures and associated audio for digital storage media at up to about 1,5 Mbit/s – Part 2: Video, 1993.

ITU-T Recommendation H.262 / ISO/IEC 13818-2 (MPEG-2 Video):
Information technology -- Generic coding of moving pictures and associated audio
information: Video, 1995.

ITU-T Recommendation H.263: Video coding for low bit rate communication, 1995; version 2, 1998; version 3, 2000.

ISO/IEC 14496-2 (MPEG-4 Visual): Information technology -- Coding of audio-visual objects – Part 2: Visual, 1999.

As such, the foregoing standards are expressly incorporated herein by reference for the purpose of illustrating certain aspects of the decoding process.

It is to be appreciated, however, that the scope of the present invention extends well beyond the particular implementations described. In describing the present invention, example network architectures and associated methods will be

described with reference to the above drawings. It is noted, however, that modification to the architecture and methods described herein may well be made without deviating from spirit and scope of the present invention. Indeed, such alternate embodiments are anticipated.

Terminology

It is to be appreciated that those skilled in the art employ various terms of art when describing certain aspects of multimedia content, the encoding and/or decoding process. While one skilled in the art is generally familiar with such terms, a brief list of terminology employed throughout the specification is provided to facilitate understanding of context and detail of the present invention.

BPP - a parameter specifying the number of bits per sample, e.g., eight (8).

component – one of three color channels {Y, Cb, Cr }.

host CPU – programmable processor which controls overall function of a computing environment (high level operations).

decoder – an aspect of a media processing system; an application typically executing on a host CPU to perform one or more video decoding functions.

accelerator – an aspect of a media processing system; a functional unit which executes computationally intensive, but high rate operations such as IDCT, MCP, display format conversion.

inverse discrete cosine transform (IDCT) – a transformation operation used as part of a video decoding process.

motion compensated prediction (MCP) – the stage of a video decoding process involving prediction of the values of a new picture using spatially-shifted areas of content from previously-decoded pictures.

media processing system – one or more elements which process (i.e., encode and/or decode) media content in accordance with a coding standard.

intra – representation of picture content without prediction using any previously-decoded picture as a reference.

inter – representation of picture content by first encoding a prediction of an area of the picture using some previously-decoded picture and then optionally adding a signal representing the deviation from that prediction.

residual difference decoding – decoding of the waveform which represents the error signal which has been encoded to represent whatever signal remains after motion-compensated prediction as appropriate. This may entail simply an “intra” representation of a non-predicted waveform or an “inter” difference after prediction.

4:2:0 sampling – a method of representing an image using twice as many luminance (Y) samples, both horizontally and vertically, relative to the number of samples used for the chrominance (Cb and Cr) components.

macroblock – a set of data comprising the samples necessary to represent a particular spatial region of picture content, including one or more blocks of all color channel components of a video signal. For example, current video coding standards often use 4:2:0 sampling with macroblocks consisting of four 8x8 blocks of Y component data and one 8x8 block of Cb and one 8x8 block of Cr data to represent each 16x16 area of picture content.

globally-unique identifier (GUID) – a 128-bit number used as a unique item identity indication.

Example Computer System

1 In the discussion herein, the invention is introduced in the general context
2 of computer-executable instructions, such as program modules, application
3 program interfaces, and the like, being executed by one or more computing
4 devices. Generally, such application program interfaces, program modules and the
5 like include routines, programs, objects, components, data structures, etc. that
6 perform particular tasks or implement particular abstract data types. Moreover,
7 those skilled in the art will appreciate that the invention may be practiced with any
8 of a number of alternate computing devices/computing configurations including,
9 for example, a personal computer, hand-held devices, personal digital assistants
10 (PDA), a KIOSK, multiprocessor systems, microprocessor-based or programmable
11 consumer electronics, network PCs, minicomputers, mainframe computers, and
12 the like. In a distributed computer environment, program modules may be located
13 in both local and remote memory storage devices. It is to be appreciated, however,
14 that the present invention may alternatively be implemented in hardware such as,
15 for example, a microcontroller, a processor, an application specific integrated
16 circuit (ASIC), a field-programmable gate array (FPGA), a programmable logic
17 device (PLD), and the like.

18 **Fig. 1** shows a general example of a computing system 102 incorporating
19 the teachings of the present invention. It will be evident, from the discussion to
20 follow, that computer 102 is intended to represent any of a class of general or
21 special purpose computing platforms which, when endowed with the innovative
22 multimedia application program interface (API) 104, implement the teachings of
23 the present invention. In this regard, the following description of computer system
24 102 is intended to be merely illustrative, as computer systems of greater or lesser
25

1 capability may well be substituted without deviating from the spirit and scope of
2 the present invention.

3 As shown, computer 102 includes one or more processors or processing
4 units 132, a system memory 134, and a bus 136 that couples various system
5 components including the system memory 134 to processors 132.

6 The bus 136 represents one or more of any of several types of bus
7 structures, including a memory bus or memory controller, a peripheral bus, an
8 accelerated graphics port (AGP), and a processor or local bus using any of a
9 variety of bus architectures. According to one implementation, a decoder
10 application executing on processing unit 132 communicates with a video
11 accelerator via the Personal Computer Interface Accelerated Graphics Port
12 (PCI/AGP) bus. The system memory includes read-only memory (ROM) 138 and
13 random access memory (RAM) 140. A basic input/output system (BIOS) 142,
14 containing the basic routines that help to transfer information between elements
15 within computer 102, such as during start-up, is stored in ROM 138. Computer
16 102 further includes a hard disk drive 144 for reading from and writing to a hard
17 disk, not shown, a magnetic disk drive 146 for reading from and writing to a
18 removable magnetic disk 148, and an optical disk drive 150 for reading from or
19 writing to a removable optical disk 152 such as a CD ROM, DVD ROM or other
20 such optical media.

21 The hard disk drive 144, magnetic disk drive 146, and optical disk drive
22 150 are connected to the bus 136 by a SCSI interface 154 or some other suitable
23 bus interface. The drives and their associated computer-readable media provide
24 nonvolatile storage of computer readable instructions, data structures, program
25 modules and other data for computer 102.

1 through an interface 170 that is coupled to bus 136. A monitor 172 or other type
2 of display device is also connected to the bus 136 via an interface, such as a video
3 adapter 174. In addition to the monitor 172, personal computers often include
4 other peripheral output devices (not shown) such as speakers and printers.

5 As shown, computer 102 operates in a networked environment using
6 logical connections to one or more remote computers, such as a remote computer
7 176. The remote computer 176 may be another personal computer, a personal
8 digital assistant, a server, a router or other network device, a network "thin-client"
9 PC, a peer device or other common network node, and typically includes many or
10 all of the elements described above relative to computer 102, although only a
11 memory storage device 178 has been illustrated in Fig. 1.

12 As shown, the logical connections depicted in Fig. 1 include a local area
13 network (LAN) 180 and a wide area network (WAN) 182. Such networking
14 environments are commonplace in offices, enterprise-wide computer networks,
15 Intranets, and the Internet. In one embodiment, remote computer 176 executes an
16 Internet Web browser program such as the "Internet Explorer" Web browser
17 manufactured and distributed by Microsoft Corporation of Redmond, Washington
18 to access and utilize online services.

19 When used in a LAN networking environment, computer 102 is connected
20 to the local network 180 through a network interface or adapter 184. When used
21 in a WAN networking environment, computer 102 typically includes a modem 186
22 or other means for establishing communications over the wide area network 182,
23 such as the Internet. The modem 186, which may be internal or external, is
24 typically connected to the bus 136 via a serial port interface 156. In a networked
25 environment, program modules depicted relative to the personal computer 102, or

Example API Architecture and Functional Relationships

Fig. 2 illustrates a block diagram of an example architecture for an adaptive multimedia API 104, as well as the functional relationships of API 104 to multimedia accelerator(s) 174 and decoder application(s) 160. According to the illustrated example embodiment, adaptive multimedia API 104 facilitates communication between a host processing unit 132, which executes one or more decoder applications (e.g., 160A-N) to render received multimedia content for a user, and one or more multimedia accelerator's 174A-N. According to one aspect of the invention, to be described more fully below, API 104 is not specific to any particular multimedia application 160A-N, host processor 132 and/or multimedia accelerator 174A-N (cumulatively referred to as a multimedia processing system). Unlike prior art multimedia API's which are designed to work with a particular media processing system, API 104 identifies the operational capability of one or more of the multimedia processing system elements and selectively negotiates the processing of received multimedia content across these elements to improve multimedia processing performance. Thus, API 104 may be utilized to facilitate the interoperability of any decoder application with any video decoder accelerator.

As introduced above, in general, an API may well comprise one or more of executable functions, messages, data structures and data types that enable an application to interface with one or more hardware devices. Thus, according to the illustrated example embodiment of Fig. 2, multimedia API 104 is comprised of one or more data structures including one or more auto-negotiation data structure(s) 202 and one or more operational data structure(s) 204.

According to one aspect of the present invention, to be described more fully below, the auto-negotiation data structure(s) 202 of API 104 are selectively invoked by a media processing system element to identify the media processing capability of the media processing system, whereupon API 104 selects one or more operational data structure(s) 204 appropriate to facilitate the negotiated processing of the media among and between the processing system elements. In this regard, API 104 facilitates the processing of media content without *a priori* knowledge of the processing capability of the elements comprising the media processing system.

Auto-negotiation Data Structure(s)

As used herein, the auto-negotiation data structure(s) 202 are a series of commands, invoked in an iterative fashion by a decoder application, for example, to identify the media decoding capability of an accelerator. According to one implementation of the present invention, the auto-negotiation data structure(s) include (1) a ConnectMode data structure, and (2) a ConnectConfig data structure. According to one implementation, the ConnectMode data structure specifies a proposed mode of operation and/or a proposed video decode format (e.g., MPEG-1, MPEG-2, etc.). A number of alternate modes of operation may well be implemented and defined within the ConnectMode data structure(s) such as, for example, an MPEG-2 mode wherein the API only invokes those data formats necessary for MPEG-2 decoding without further negotiation of other data formats, a protected mode (i.e., utilizing encrypted communication between the decoder and the accelerator), or a normal mode (i.e., non-restricted, non-protected).

The ConnectConfig data structure provides information on how the API 104 is to be configured to decode the video in accordance with the video format

1 identified in the ConnectMode data structure. According to one illustrative
2 example, the ConnectConfig data structure includes information regarding
3 intermediate data formats to be used (if any), which aspects of the decoding
4 process will reside on the host versus the accelerator, and the like. According to
5 one embodiment, the ConnectMode and ConnectConfig data structures are
6 iteratively passed between the decoder and the accelerator utilizing a ConnectInfo
7 command, e.g., ConnectInfo {ConnectMode, ConnectConfig}. The ConnectMode
8 and ConnectConfig data structures can be looked upon as two "orthogonal"
9 aspects of codec construction between the decoder software and video accelerator
10 driver.

11 According to one implementation, decoder 160 issues the ConnectInfo
12 command with one of a number of ConnectMode and ConnectConfig
13 combinations, to accommodate any of a number of multimedia codecs. If the
14 accelerator 174 does not support a particular ConnectMode/ConnectConfig
15 combination, a negative response to the ConnectInfo command is sent to the
16 decoder 160. If, however, the accelerator 174 does support the Mode/Config
17 combination, a positive response is issued to decoder 160, as API 104 selects
18 appropriate ones of the operational data structure(s) 204 to facilitate the decoding
19 of the multimedia in the mutually agreed upon format. According to one
20 implementation, API 104 selects a ConnectMode/ConnectConfig combination
21 reflecting the MPEG-2 main profile, main level with host-based IDCT as a default
22 proposal, followed by other combinations. Example ConnectMode and
23 ConnectConfig parameters are introduced with reference to Table I and Table II,
24 respectively, below.

25 ConnectMode {

```

ModeGUID      (128b; The Global ID of the Intended Mode)
dwRestrictedMode (16b; Restricted Mode ID)
}

```

Table I: Example ConnectMode Data Structure Settings

As introduced in Table I, above, the ConnectMode data structure passes the GUID of a proposed mode of operation. In addition, in accordance with the illustrated example embodiment, a restricted mode may also be negotiated within the ConnectMode data structure.

```

ConnectConfig {
    //Encryption GUIDs
    ConfigBitstreamEncryptionGUID
    ConfigMBcontrolEncryptionGUID
    ConfigResidDiffEncryptionGUID
    //Bitstream Processing Indicator
    ConfigBitstreamRaw
    //Macroblock Control Configuration
    ConfigMBcontrolRaasterOrder
    //Host Residual Difference Configuration
    ConfigResidDiffHost
    ConfigSpatialResid8
    ConfigOverflowBlocks
    ConfigResid8Subtraction
    ConfigSpatialHost8or9Clipping
    //Accelerator Residual Difference Configuration
    ConfigResidDiffAccelerator
    ConfigHostInverseScan
    ConfigSpecificIDCT
}

```

Table II: Example ConnectConfig Data Structure Parameters

With reference to Table II, a number of operational parameters are negotiated within the ConnectConfig data structure including, but not limited to, encryption parameters, bitstream processing indicator, macroblock control configuration information, host residual difference configuration information and accelerator residual difference configuration information. An example implementation of each of the ConnectConfig parameters are introduced, below.

shall be in raster-scan order, and a value of “0” indicates arbitrary order. For some types of bitstreams, forcing raster order will either greatly increase the number of required macroblock control buffers that must be processed or will require host reordering of the control information. Support of arbitrary order can thus be advantageous for the decoding process. For example, H.261 CIF-resolution decoding can require 36 macroblock control buffers per picture if raster-scan order is necessary within each buffer (H.263 Annex K’s arbitrary slice ordering and rectangular slice modes have similar repercussions.) An intermediate-term requirement is to support “0”. Additional support of “1” is desired.

ConfigResidDiffHost: A value of “1” specifies that some residual difference decoding data may be sent as blocks in the spatial domain from the host, and a value of “0” specifies that spatial domain data will not be sent. Shall be “0” if ConfigBitstreamRaw is “1”. An intermediate-term requirement is to support “1”, which is the preferred value.

ConfigSpatialResid8: A value of “1” indicates that host residual difference spatial-domain blocks of prediction residual data for predicted pictures will be sent using 8 bit signed samples, and a value of “0” indicates that such blocks are sent using 16 bit signed samples. (For intra macroblocks, these signed samples are sent relative to a constant reference value of 2^{BPP-1} .) Shall be “0” if ConfigResidDiffHost is “0”.

ConfigOverflowBlocks: A value of “1” indicates that host residual difference spatial blocks of prediction residual data for predicted pictures may be sent using 8 bit signed “overflow” blocks in a second pass for each macroblock rather than sending only one set of signed block data, and a value of “0” indicates that such overflow blocks shall not be sent (instead using a second complete pass

for any necessary overflow blocks, such as a “read-modify-write” picture as described below). Shall be “0” if ConfigSpatialResid8 is “0”. When ConfigSpatialResid8 is “1”, a value of “1” for ConfigOverflowBlocks is considered preferred over a value of “0”, as it prevents the need for two complete macroblock control command passes to create a single output picture. An intermediate-term requirement is support of “1” if ConfigSpatialResid8 = “1” is supported.

ConfigResid8Subtraction: A value of “1” when ConfigSpatialResid8 is “1” indicates that 8-bit differences can be subtracted rather than added. Shall be “0” unless ConfigSpatialResid8 is “1”. If “1” with ConfigOverflowBlocks equal to “1”, this indicates that any overflow blocks will be subtracted rather than added. If “1” with ConfigOverflowBlocks equal to “0”, this indicates that frames may be sent with single-pass subtracted 8-bit spatial differences. An intermediate-term requirement is to support “1” if ConfigSpatialResid8 is “1”.

ConfigSpatialHost8or9Clipping: A value of “1” indicates that spatial-domain intra blocks shall be clipped to an 8-bit range on the host and that spatial-domain inter blocks shall be clipped to a 9-bit range on the host, and a value of “0” indicates that any necessary clipping is performed on the accelerator. An intermediate-term requirement is to support “0”. Nearer-term support of “1” is allowed but less preferred, and is considered a lower level of accelerator capability.

ConfigSpatialResidInterleaved: A value of “1” when ConfigResidDiffHost is “1” and the YUV format is “NV12” or “NV21” indicates that any spatial-domain residual difference data shall be sent in a chroma-interleaved form matching the YUV format chroma interleaving pattern. Shall be

“0” unless ConfigResidDiffHost is “1” and the YUV format is “NV12” or “NV21”. An intermediate-term requirement is to support “0”. Nearer-term support of “1” is allowed but less preferred, and is considered a lower level of accelerator capability.

ConfigResidDiffAccelerator: A value of “1” indicates that transform-domain blocks of coefficient data may be sent from the host for accelerator-based IDCT, and a value of “0” specifies that accelerator-based IDCT will not be used. If both ConfigResidDiffHost and ConfigResidDiffAccelerator are “1”, this indicates that some residual difference decoding will be done on the host and some on the accelerator, as indicated by macroblock-level control commands. Shall be “0” if ConfigBitstreamRaw is “1”. Support for ConfigResidDiffAccelerator equal to “1” is desired, but there is not expected to be an intermediate-term requirement for this support. Support for ConfigResidDiffAccelerator being “1” with ConfigResidDiffHost also being “1” indicates that the residual difference decoding can be shared between the host and accelerator on a macroblock basis, and is considered an even higher level of accelerator capability than ConfigResidDiffAccelerator being “1” with ConfigResidDiffHost being “0”.

ConfigHostInverseScan: A value of “1” indicates that the inverse scan for transform-domain block processing will be performed on the host, and absolute indices will be sent instead for any transform coefficients, and a value of “0” indicates that inverse scan will be performed on the accelerator. Shall be “0” if ConfigResidDiffAccelerator is “0”. An intermediate-term expected requirement is to support “1” if ConfigResidDiffAccelerator is “1”. Nearer-term support of “0” is allowed but less preferred, and is considered a lower level of accelerator capability.

1 **ConfigSpecificIDCT:** A value of “1” indicates use of the IDCT specified
2 in ITU-T H.263 Annex W, and a value of “0” indicates that any compliant IDCT
3 can be used for off-host IDCT. Shall be zero if ConfigResidDiffAccelerator is “0”
4 (indicating purely host-based residual difference decoding). An intermediate-term
5 expected requirement is to support “0” if ConfigResidDiffAccelerator is “1”.
6 Additional support of “1” is desired and is considered a higher level of accelerator
7 capability.

8 9 **Operational Data Structure(s)**

10 In addition to the auto-negotiation data structure(s) 202, API 104 also
11 includes one or more operational data structure(s) 204. As introduced above, one
12 or more of the operational data structure(s) 204 are selectively invoked by API
13 104 to facilitate the communication required to effect the negotiated division in
14 media decoding among and between media processing system elements (e.g.,
15 decoder application and accelerator). In accordance with the illustrated example
16 embodiment of a video decoding system, the operational data structure(s) 204
17 include picture level parameters and/or buffer structure for macroblocks of a
18 picture. The picture level parameters the buffer structure required for media
19 decoding depends, at least in part, on which elements of the media processing
20 system will be to perform the various decoding tasks. According to one
21 implementation, API 104 facilitates configuration of a number of picture level
22 parameter(s) (see, e.g., Table III below), and dynamically adapts buffer
23 structure(s) to accommodate Pre-IDCT saturation, Mismatch Control, IDCT,
24 Picture Reconstruction, and Reconstruction Clipping (each of which are discussed
25 in turn, below).

Picture-Level Parameters

One or more picture level parameters are sent using a PictureParameters{} command within the operational data structure 204 defining a number of picture-level variables once per picture between decoder application and the accelerator. In accordance with the illustrated example embodiment, the picture level parameters of the operational data structure describe one or more aspects of the picture to be decoded such as, for example, one or more picture indices (e.g., decoded picture index, deblocked picture index, etc.), the picture encoding type (e.g., intra-encoded, inter-encoded, etc.), and the like. An example of set of picture level parameters are provided with reference to Table III, below.

```
PictureParameters {
    DecodedPictureIndex
    DeblockedPictureIndex
    SubpictureBlendedIndex
    ForwardRefPictureIndex
    BackwardRefPictureIndex

    IntraPicture
    BPPminus1
    SecondField
    SubpictureControlPresent
    ReservedBits

    MacroblockWidthMinus1
    MacroblockHeightMinus1

    BlockWidthMinus1
    BlockHeightMinus1

    PicWidthInMinus1
    BlockHeightInMinus1

    ChromaFormat
    PicStructure
    Rcontrol
    BidirectionalAveragingMode
    MVprecisionAndChromaRelation

    ReservedBits
    PicSpatialResid8
    PicOverflowBlocks
    PicResid8Subtraction
    PicExtrapolation
}
```

```

PicDeblocked

Pic4Mvallowed
PicOBMC
PicBinPB
MV_RPS
PicDeblockedConfined
PicReadbackRequests
ReservedBits

PicScanFixed
PicScanMethod
Reserved Bits

PicReampleOn
PicResampleBefore
PicResampleRcontrol
ReservedBits

PicResampleSourcePicIndex
PicResampleDestPicIndex

PicResampleSourceWidthMinus1
PicResampleSourceHeightMinus1

PicResampleDestWidthMinus1
PicResampleDestHeightMinus1

PicResampleFullDestWidthMinus1
PicResampleFullDestHeightMinus1
}

```

Table III: Example Picture-level Parameters

In accordance with one example implementation, each of the foregoing parameters will be defined, in turn, below:

DecodedPictureIndex: Specifies destination frame buffer for the decoded macroblocks.

DeblockedPictureIndex: Specifies destination frame buffer for the deblocked output picture when $bPicDeblocked = 1$. Has no meaning and shall be zero if $bPicDeblocked = 0$. May be the same as $wDecodedPictureIndex$.

SubpictureBlendedIndex: Specifies destination frame buffer for the output picture after blending with a DVD subpicture. Subpicture blending shall occur after deblocking if applicable. Shall be equal to $wDeblockedPictureIndex$ or

wDecodedPictureIndex as applicable if no subpicture blending is required for the picture.

ForwardRefPictureIndex: Specifies the frame buffer index of the picture to be used as a reference picture for “forward prediction” of the current picture. Shall not be the same as DecodedPictureIndex unless all motion prediction for the current picture uses forward motion with zero-valued motion vectors and no macroblocks are sent as intra and PicSpatialResid8 is 1 and PicOverflowBlocks is 0 and PicResid8Subtraction is 1. NOTE: The ability for wForwardRefPictureIndex to be set equal to wDecodedPictureIndex if all motion prediction uses forward prediction with zero-valued motion vectors is provided to allow processing of 8-bit difference pictures (see PicSpatialResid8, PicOverflowBlocks, and PicResid8Subtraction below) by a two-picture pass process – one pass of decoding to perform motion compensation and to add the first set of 8-bit differences, and a second pass to perform “read-modify-write” operations to subtract a second set of 8-bit differences and obtain the final result.

BackwardRefPictureIndex: Specifies the frame buffer index of the picture to be used as a reference picture for “backward prediction” of the current picture. Shall not be the same as DecodedPictureIndex if backward reference motion prediction is used.

IntraPicture: Indicates whether motion prediction is needed for this picture. If IntraPicture = 1, no motion prediction is performed for the picture. Otherwise, motion prediction information shall be sent for the picture.

BPPminus1: Specifies the number of bits per pixel for the video sample values. This shall be at least 7. It is equal to 7 for MPEG-1, MPEG-2, H.261, and

H.263. A larger number of bits per pixel is supported in some operational modes of MPEG-4. A derived term called **BPP** is formed by adding one to **bBPPminus1**.

SecondField: Indicates whether, in the case of field-structured motion prediction, the current field is the second field of a picture. This is used to determine whether motion compensation prediction is performed using the reference picture or the opposite-parity field of the current picture.

SubpictureControlPresent: Indicates whether a subpicture control buffer is sent for the current picture.

MacroblockWidthMinus1: Specifies the destination luminance sample width of a macroblock. This is equal to 15 for MPEG-1, MPEG-2, H.263, and MPEG-4. A derived term called **MacroblockWidth** is formed by adding one to **MacroblockWidthMinus1**.

MacroblockHeightMinus1: Specifies the destination luminance sample height of a macroblock. This is equal to 15 for MPEG-1, MPEG-2, H.261, H.263, and MPEG-4. A derived term called **MacroblockHeight** is formed by adding one to **MacroblockHeightMinus1**.

BlockWidthMinus1: Specifies the block width of an residual difference block. This is equal to 7 for MPEG-1, MPEG-2, H.261, H.263, and MPEG-4. Residual difference blocks within a macroblock are sent in the order specified as in H.262 Figures 6-10, 6-11, and 6-12 (raster-scan order for Y, followed by all 4:2:0 blocks of Cb in raster-scan order, followed by 4:2:0 blocks of Cr, followed by 4:2:2 blocks of Cb, followed by 4:2:2 blocks of Cr, followed by 4:4:4 blocks of Cb, followed by 4:4:4 blocks of Cr). A derived term called W_T is formed by adding one to **BlockWidthMinus1**.

BlockHeightMinus1: Specifies the block height of an IDCT block. This is equal to 7 for MPEG-1, MPEG-2, H.261, H.263, and MPEG-4. A derived term called H_T is formed by adding one to BlockHeightMinus1.

PicWidthInMBminus1: Specifies the width of the current picture in units of macroblocks, minus 1. A derived term called **PicWidthInMB** is formed by adding one to PicWidthInMBminus1.

PicHeightInMBminus1: Specifies the width of the current picture in units of macroblocks, minus 1. A derived term called **PicHeightInMB** is formed by adding one to PicHeightInMBminus1.

ChromaFormat: Affects number of prediction error blocks expected by the Accelerator. This variable is defined in Section 6.3.5 and Table 6-5 of H.262. For MPEG-1, MPEG-2 “Main Profile,” H.261 and H.263 bitstreams, this value shall always be set to ‘01’, indicating “4:2:0” format. If ‘10’ this indicates “4:2:2”, and “11” indicates “4:4:4” sampling. Horizontal chroma siting differs slightly between H.261, H.263, MPEG-1 versus MPEG-2 and MPEG-4. This difference may be small enough to ignore.

PicStructure: This parameter has the same meaning as the *picture_structure* parameter defined in Section 6.3.10 and Table 6-14 of MPEG-2, and indicates whether the current picture is a top-field picture (value ‘01’), a bottom-field picture (value ‘10’), or a frame picture (value ‘11’). In progressive-scan frame-structured coding such as in H.261, PicStructure shall be ‘11’.

RCONTROL: This flag is defined in H.263 Section 6.1.2. It defines the rounding method to be used for half-sample motion compensation. A value of 0 indicates the half-sample rounding method found in MPEG-1, MPEG-2, and the first version of H.263. A value of 1 indicates the rounding method which includes

and a value of 0 indicates that they cannot. Shall be 0 if ConfigResidDiffHost is 0 or if BPP > 8. Shall be 1 if BPP = 8 and IntraPicture = 1 and ConfigResidDiffHost is "1". If 1, this indicates that spatial-domain intra macroblocks are sent as signed 8-bit difference values relative to the constant value 2^{BPP-1} and that spatial-domain non-intra macroblock differences are sent as signed 8-bit difference values relative to some motion compensated prediction. PicSpatialResid8 differs from ConfigSpatialResid8 in that it is an indication for a particular picture, not a global indication for the entire video sequence. In some cases such as in an intra picture with BPP equal to "8", PicSpatialResid8 will be 1 even though ConfigSpatialResid8 may be 0.

PicOverflowBlocks: A value of 1 indicates that spatial-domain difference blocks for host-based residual difference decoding can be sent using "overflow" blocks, and a value of 0 indicates that they cannot. Shall be 0 if ConfigResidDiffHost is 0 or if BPP > 8. PicOverflowBlocks differs from ConfigOverflowBlocks in that it is an indication for a particular picture, not a global indication for the entire video sequence. In some cases such as in an intra picture with BPP equal to "8", PicOverflowBlocks will be 0 even though ConfigOverflowBlocks is "1".

PicResid8Subtraction: A value of 1 when PicSpatialResid8 is 1 indicates that some 8-bit spatial-domain residual differences shall be subtracted rather than added, according to one aspect of the present invention. Shall be 0 if PicSpatialResid8 is 0 or ConfigResid8Subtraction is 0. According to one aspect of the present invention, if PicResid8Subtraction is 1 and PicOverflowBlocks is 1, this indicates that the spatial-domain residual difference overflow blocks shall be subtracted rather than added. If PicResid8Subtraction is 1 and PicOverflowBlocks

is 0, this indicates that no overflow blocks are sent and that all spatial-domain residual difference blocks shall be subtracted rather than added, and that no macroblocks will be sent as intra macroblocks. This ability to subtract differences rather than add them allows 8-bit difference decoding to be fully compliant with the full ± 255 range of values required in video decoder specifications, since +255 cannot be represented as the addition of two signed 8-bit numbers but any number in the range ± 255 can be represented as the difference between two signed 8-bit numbers ($+255 = +127$ minus -128). In this regard, API 104 provides a flexible solution to host-based IDCT.

PicExtrapolation: This flag indicates whether motion vectors over picture boundaries are allowed as specified by H.263 Annex D and MPEG-4. This requires either allocation of picture planes which are two macroblocks wider (one extra macroblock at the left and another at the right) and two macroblocks taller (one extra macroblock at the top and another at the bottom) than the decoded picture size, or clipping of the address of each individual pixel access to within the picture boundaries. Macroblock addresses in this specification are for macroblocks in the interior of the picture, not including padding.

PicDeblocked: Indicates whether deblocking commands are sent for this picture for creating a deblocked output picture in the picture buffer indicated in DeblockedPictureIndex. If PicDeblocked = 1, deblocking commands are sent and the deblocked frame shall be generated, and if PicDeblocked = 0, no deblocking commands are sent and no deblocked picture shall be generated.

Pic4MVallowed: Specifies whether four forward-reference motion vectors per macroblock are allowed as used in H.263 Annexes F and J.

1 **PicOBMC:** Specifies whether motion compensation for the current picture
2 operates using overlapped block motion compensation (OBMC) as specified in
3 H.263 Annex F. Shall be zero if Pic4MVallowed is 0.

4 **PicBinPB:** Specifies whether bi-directionally-predicted macroblocks in the
5 picture use “B in PB” motion compensation, which restricts the bi-directionally
6 predicted area for each macroblock to the region of the corresponding macroblock
7 in the backward reference picture, as specified in Annexes G and M of H.263.

8 **MV_RPS:** Specifies use of motion vector reference picture selection. If 1,
9 this indicates that a reference picture index is sent for each motion vector rather
10 than just forward and possibly backward motion picture indexes for the picture as
11 a whole. If MV_RPS is 1, the parameters ForwardRefPictureIndex and
12 BackwardRefPictureIndex have no meaning and shall be zero.

13 **PicDeblockConfined:** Indicates whether deblocking filter command
14 buffers contain commands which confine the effect of the deblocking filter
15 operations to within the same set of macroblocks as are contained in the buffer.

16 **PicReadbackRequests:** Indicates whether read-back control requests are
17 issued for the current picture to read back the values of macroblocks in the final
18 decoded picture. A value of 1 indicates that read-back requests are present, and 0
19 indicates that they are not.

20 **PicScanFixed:** When using accelerator-based IDCT processing of residual
21 difference blocks, a value of 1 for this flag indicates that the inverse-scan method
22 is the same for all macroblocks in the picture, and a value of 0 indicates that it is
23 not. Shall be 1 if ConfigHostInverseScan is 1 or if ConfigResidDiffAccelerator is
24 0.

PicScanMethod: When PicScanFixed is 1, this field indicates the fixed inverse scan method for the picture. When PicScanFixed is 0, this field has no meaning and shall be '00'. If PicScanFixed = 1 this field shall have one of the following values:

If ConfigHostInverseScan = 0, PicScanMethod shall be as follows:

'00' = Zig-zag scan (H.262 Figure 7-2),

'01' = Alternate-vertical (H.262 Figure 7-3),

'10' = Alternate-horizontal (H.263 Figure I.2 Part a),

If ConfigHostInverseScan = 1, PicScanMethod shall be as follows:

'11' = Arbitrary scan with absolute coefficient address.

PicResampleOn: Specifies whether an input picture is to be resampled to a destination buffer prior to decoding the current picture or whether the final output picture is to be resampled for use as an upsampled display picture or as a future upsampled or downsampled reference picture. The resampling is performed as specified for H.263 Annex O Spatial Scalability or for H.263 Annex P, which we believe to be the same as in some forms of the Spatial Scalability in MPEG-2 and MPEG-4. If this value is 1, the remaining resampling parameters are used to control the resampling operation. If 0, the resampling is not performed and the remaining resampling parameters shall be zero. If PicExtrapolation is 1 and the padding method is used on the accelerator, any resampling shall include padding of the resampled picture as well – and this padding shall be at least one macroblock in width and height around each edge of the resampled picture regardless of the resampling operation which is performed.

PicResampleBefore: Specifies whether the resampling process is to be applied before (a value of 1) the processing of the current picture, or after it (a

value of 0). If resampling after decoding is indicated and DeblockedPictureIndex differs from DecodedPictureIndex, the decoded picture (not the deblocked picture) is the one that has the resampling applied to it. If resampling after decoding is indicated and the DeblockedPictureIndex is the same as the DecodedPictureIndex, the deblocking shall be applied to the decoded picture with the result placed in that same destination frame buffer – and the resampling process shall be performed using the deblocked frame buffer as the input picture.

PicResampleRcontrol: Specifies the averaging rounding mode of the resampling operation. In the case of H.263 Annex O Spatial Scalability, this parameter shall be 1. (This corresponds to the value of *RCRPR* in H.263 Annex P which is equivalent to the upsampling needed for H.263 Annex O spatial scalability.) In the case of H.263 Annex P Reference Picture Resampling, this parameter shall be equal to the H.263 parameter *RCRPR*.

PicResampleSourcePicIndex: Specifies the reference buffer to be resampled in order to make it the same size as the current picture.

PicResampleDestPicIndex: Specifies the buffer to be used for the output of the reference picture resampling operation. This buffer can then be used as a reference picture for decoding the current picture.

PicResampleSourceWidthMinus1: Specifies the width of the area of the source picture to be resampled to the destination picture. A derived parameter *PicResampleSourceWidth* is formed by adding one to *PicResampleSourceWidth*.

PicResampleSourceHeightMinus1: Specifies the height of the area of the source picture to be resampled to the destination picture. A derived parameter *PicResampleSourceHeight* is formed by adding one to *PicResampleSourceHeight*.

PicResampleDestWidthMinus1: Specifies the width of the area of the destination picture to contain the resampled data from the source picture. A derived parameter `PicResampleDestWidth` is formed by adding one to `PicResampleDestWidth`.

PicResampleDestHeightMinus1: Specifies the height of the area of the destination picture to contain the resampled data from the source picture. A derived parameter `PicResampleDestHeight` is formed by adding one to `PicResampleSourceHeight`.

PicResampleFullDestWidthMinus1: Specifies the full height of the area of the destination picture to contain the resampled data from the source picture. Clipping shall be used to generate any samples outside the source resampling area. (This parameter is necessary for H.263 Annex P support of custom source formats in which the luminance width is not divisible by 16.) A derived parameter **PicResampleFullDestWidth** is formed by adding one to **PicResampleFullDestWidth**.

PicResampleFullDestHeightMinus1: Specifies the full height of the area of the destination picture to contain the resampled data from the source picture. Clipping shall be used to generate any samples outside the source resampling area. (This parameter is necessary for H.263 Annex P support of custom source formats in which the luminance height is not divisible by 16.) A derived parameter **PicResampleFullDestWidth** is formed by adding one to **PicResampleFullDestHeight**.

Buffer Structure for Macroblocks of a Picture

1 The processing of the picture requires that motion prediction for each
2 macroblock must precede the addition of the IDCT residual data. According to
3 one implementation of the present invention, this is accomplished either by
4 processing the motion prediction commands first and then reading this data back
5 in from the destination picture buffer while processing the IDCT residual coding
6 commands, or by processing these two buffers in a coordinated fashion, i.e.,
7 adding the residual data to the prediction before writing the result to the
8 destination picture buffer. The motion prediction command and IDCT residual
9 coding command for each macroblock affect only the rectangular region within
10 that macroblock.

11 A deblocking filter command for a macroblock may require access to read
12 the reconstructed values of two rows and two columns of samples neighboring the
13 current macroblock at the top and left as well as reconstructed values within the
14 current macroblock. It can result in modification of one row and one column of
15 samples neighboring the current macroblock at the top and left as well as three
16 rows and three columns within the current macroblock. The filtering process for a
17 given macroblock may therefore require the prior reconstruction of other
18 macroblocks. Two different types of deblocking filter buffers are defined herein:
19 (1) a buffer type which requires access and modification of the value of
20 reconstructed samples for macroblocks outside the current buffer (e.g., when
21 PicDeblockConfined is set to '0'), and (2) a buffer type which does not (e.g., when
22 PicDeblockConfined is set to '1'). To process the first of these two types of
23 deblocking command buffer, the accelerator must ensure that the reconstruction
24 has been completed for all buffers which affect macroblocks to the left and top of
25 the macroblocks in the current buffer before processing the deblocking commands

in the current buffer. Processing the second of these two types requires only prior reconstruction values within the current buffer. The deblocking post-processing can be conducted either by processing the motion prediction and IDCT residual coding commands for the entire buffer or frame first, followed by reading back in the values of some of the samples and modifying them as a result of the deblocking filter operations, or by processing the deblocking command buffer in a coordinated fashion with the IDCT residual coding buffer – performing the deblocking before writing the final output values to the destination picture buffer. Note also that the destination picture buffer for the deblocked picture may differ from that of the reconstructed picture prior to deblocking, in order to support “outside the loop” deblocking as a post-processing operation which does not affect the sample values used for prediction of the next picture.

Table IV, below, provides example macroblock control commands, selectively invoked by API 104 in operational data structure(s) 204 in response to a negotiated decoding format and media processing task allocation among and between media processing system elements.

```

if (IntraPicture)
    NumMV = 0;
else if (PicOBMC) {
    NumMV = 10;
    if (PicBinPB)
        NumMV++;
} else {
    NumMV = 4;
    if (PicBinPB && Pic4MVallowed)
        NumMV++;
}

if (ChromaFormat == '01')
    NumBlocksPerMB = 6
else if (ChromaFormat == '10')
    NumBlocksPerMB = 8
else
    NumBlocksPerMB = 12

```

```

1  MB_Control {
   // General Macroblock Info
   MBaddress
   MBtype
   MBskipsFollowing

2

3  // Residual Difference Info
   MBdataLocation
   PatternCode

4

5      if(PicOverflowBlocks==1 && IntraMacroblock==0){
        PC_Overflow
        ReservedBits2
6      } else if(HostResidDiff)
        ReservedBits3
7      else
        for(i=0; i<NumBlocksPerMB; i++)
8          NumCoef[i]

9  // Motion Prediction Info
        for(i=0; i<NumMV; i++) {
10         MVector[i].horz
           MVector[i].vert
        }
11     if(MV_RPS)
        for(i=0; i<NumMV; i++)
12         RefPicSelect[i]
        ReservedBits4
13 }

```

Table IV: Example Control Commands

Each of the various control command attributes are described, in turn, below.

MBaddress: Specifies the macroblock address of the current macroblock in raster scan order (0 being the address of the top left macroblock, PicWidthInMBminus1 being the address of the top right macroblock, and PicHeightInMBminus1 * PicWidthInMB being the address of the bottom left macroblock, and PicHeightInMBminus1 * PicWidthInMB + PicWidthInMBminus1 being the address of the bottom right macroblock).

MBtype: Specifies the type of macroblock being processed as described below:

bit 15: MvertFieldSel[3] (The MSB),

bit 14: MvertFieldSel[2],

bit 13: MvertFieldSel[1],

bit 12: MvertFieldSel[0]: Specifies vertical field selection for corresponding motion vectors sent later in the macroblock control command, as specified in further detail below. For frame-based motion with a frame picture structure (e.g., for H.261 and H.263), these bits shall all be zero. The use of these bits is the same as that specified for the corresponding bits in Section 6.3.17.2 of H.262.

bit 11: ReservedBits.

bit 10: HostResidDiff: Specifies whether spatial-domain residual difference decoded blocks are sent or whether transform coefficients are sent for off-host IDCT for the current macroblock.

bits 9 and 8: MotionType: Specifies the motion type in the picture, as specified in further detail below. For frame-based motion with a frame picture structure (e.g., for H.261 and H.263), these bits shall be equal to '10'. The use of these bits is the same as that specified for the corresponding bits in Section 6.3.17.1 and Table 6-17 of H.262.

bits 7 and 6: MBscanMethod: Shall equal PicScanMethod if PicScanFixed is 1.

If ConfigHostInverseScan = 0, MBscanMethod shall be as follows:

‘00’ = Zig-zag scan (H.262 Figure 7-2),

‘01’ = Alternate-vertical (H.262 Figure 7-3),

‘10’ = Alternate-horizontal (H.263 Figure I.2 Part a),

If ConfigHostInverseScan = 1, MBscanMethod shall be equal to:

‘11’ = Arbitrary scan with absolute coefficient address.

bit 5: FieldResidual: A flag indicating whether the IDCT blocks use a field IDCT structure as specified in H.262.

bit 4: H261LoopFilter: A flag specifying whether the H.261 loop filter (Section 3.2.3 of H.261) is active for the current macroblock prediction. The H.261 loop filter is a separable $\frac{1}{4}$, $\frac{1}{2}$, $\frac{1}{4}$ filter applied both horizontally and vertically to all six blocks in an H.261 macroblock except at block edges where one of the taps would fall outside the block. In such cases the filter is changed to have coefficients 0, 1, 0. Full arithmetic precision is retained with rounding to 8-bit integers at the output of the 2-D filter process (half-integer or higher values being rounded up).

bit 3: Motion4MV: A flag indicating that forward motion uses a distinct motion vector for each of the four luminance blocks in the macroblock, as used in H.263 Annexes F and J. Motion4MV shall be 0 if MotionForward is 0 or Pic4MVallowed is 0.

bit 2: MotionBackward: A flag used as specified for the corresponding parameter in H.262. Further information on the use of this flag is given below.

bit 1: MotionForward: A flag used as specified for the corresponding flag in H.262. Further information on the use of this flag is given below.

bit 0: IntraMacroblock: (The LSB) A flag indicating that the macroblock is coded as “intra”, and no motion vectors are used for the current macroblock. Further information on the use of this flag is given below.

MBskipsFollowing: Specifies the number of “skipped macroblocks” to be generated following the current macroblock. Skipped macroblocks shall be generated using the rules specified in H.262 Section 7.6.6. According to one implementation, the API 104 operates by using an indication of the number of skipped macroblocks *after* the current macroblock instead of the number of skipped macroblocks before the current macroblock. Insofar as the method of generating skipped macroblocks as specified in H.262 Section 7.6.6 depends on the parameters of the macroblock preceding the skipped macroblocks, specifying the operation in this way means that only the content of a single macroblock control structure need be accessed for the generation of the skipped macroblocks.

For implementation of standard video codecs other than H.262 (MPEG-2), some “skipped” macroblocks may need to be generated with some indication other than the skipped macroblock handling used by MBskipsFollowing if the skipped macroblock handling differs from that of H.262.

The generation of macroblocks indicated as skipped in H.263 with Advanced Prediction mode active requires coding some “skipped” macroblocks as non-skipped macroblocks using this specification – in order to specify the OBMC effect within these macroblocks.

MBdataLocation: An index into the IDCT residual coding block data buffer, indicating the location of the residual difference data for the blocks of the current macroblock, expressed as a multiple of 32 bits.

PatternCode: When using host-based residual difference decoding, bit $11-i$ of $wPatternCode$ (where bit 0 is the LSB) indicates whether a residual difference block is sent for block i , where i is the index of the block within the macroblock as specified in Figures 6-10, 6-11, and 6-12 (raster-scan order for Y, followed by 4:2:0 blocks of Cb in raster-scan order, followed by 4:2:0 blocks of Cr, followed by 4:2:2 blocks of Cb, followed by 4:2:2 blocks of Cr, followed by 4:4:4 blocks of Cb, followed by 4:4:4 blocks of Cr). The data for the coded blocks (those blocks having bit $11-i$ equal to 1) is found in the residual coding buffer in the same indexing order (increasing i). For 4:2:0 H.262 data, the value of $wPatternCode$ corresponds to shifting the decoded value of CBP left by six bit positions (those lower bit positions being for the use of 4:2:2 and 4:4:4 chroma formats).

If $ConfigSpatialResidInterleaved$ is “1”, host-based residual differences are sent in a chroma-interleaved form matching that of the YUV pixel format in use. In this case each Cb and spatially-corresponding Cr pair of blocks is treated as a single residual difference data structure unit. This does not alter the value or meaning of $PatternCode$, but it implies that both members of each pair of Cb and Cr data blocks are sent whenever either of these data blocks has the corresponding bit set in $PatternCode$. If the bit in $PatternCode$ for a particular data block is zero,

1 the corresponding residual difference data values shall be sent as zero whenever
2 this pairing necessitates sending a residual difference data block for a block with a
3 PatternCode bit equal to zero.

4 **PC_Overflow:** When using host-based residual difference decoding with
5 PicOverflowBlocks (the innovative 8-8 overflow method introduced above, and
6 described in greater detail below), PC_Overflow contains the pattern code of the
7 overflow blocks as specified in the same manner as for PatternCode. The data for
8 the coded overflow blocks (those blocks having bit 11- i equal to 1) is found in the
9 residual coding buffer in the same indexing order (increasing i).

10 **NumCoef[i]:** Indicates the number of coefficients in the residual coding
11 block data buffer for each block i of the macroblock, where i is the index of the
12 block within the macroblock as specified in H.262 Figures 6-10, 6-11, and 6-12
13 (raster-scan order for Y, followed by 4:2:0 blocks of Cb in raster-scan order,
14 followed by 4:2:0 blocks of Cr, followed by 4:2:2 blocks of Cb, followed by 4:2:2
15 blocks of Cr, followed by 4:4:4 blocks of Cb, followed by 4:4:4 blocks of Cr).
16 The data for these coefficients is found in the residual difference buffer in the
17 same order.

18 **MVector[i].horz, MVector[i].vert:** Specifies the value of a motion vector
19 in horizontal and vertical dimensions. The two-dimensional union of these two
20 values is referred to as MVvalue[i]. Each dimension of each motion vector
21 contains a signed integer motion offset in half-sample units. Both elements shall
22 be even if MVprecisionAndChromaRelation = '10' (H.261-style motion
23 supporting only integer-sample offsets).

24 **RefPicSelect[i]:** Specifies the reference picture buffer used in prediction
25 for MVvalue[i] when motion vector reference picture selection is in use.

IDCT Support

According to one aspect of the present invention, API 104 supports at least three (3) low-level methods of handling inverse discrete cosine transform (IDCT) decoding via the operational data structure(s) 204. In all cases, the basic inverse quantization process, pre-IDCT range saturation, and mismatch control (if necessary) is performed by the decoder 160 (e.g., on the host), while the final picture reconstruction and reconstruction clipping is done on the accelerator 174. The first method is to pass macroblocks of transform coefficients to the accelerator 174 for external IDCT, picture reconstruction, and reconstruction clipping. The second and third methods involve performing an IDCT by the decoder 160 and passing blocks of spatial-domain results for external picture reconstruction and clipping on the accelerator 174.

According to one implementation (also denoted with reference to Fig. 6), the pre-IDCT saturation, mismatch control, IDCT, picture reconstruction and clipping processes are defined as:

- (1) Saturating each reconstructed coefficient value in the transform coefficient block to the allowable range (typically performed by the decoder 160):

$$-2^{\text{BPP} + \log_2 \sqrt{W_T H_T}} \leq F'(u, v) \leq 2^{\text{BPP} + \log_2 \sqrt{W_T H_T}} - 1 \quad (1)$$

- (2) Mismatch control (as necessary in association with MPEG-2 decoding) is performed by adding the saturated values of all coefficients in the macroblock. According to one implementation, this is performed by XORing the least significant bits. If the sum is even, then the saturated value of the last coefficient $F'(W_T-1, H_T-1)$ is modified by subtracting one if it is odd, or adding one if it is even. The coefficient values

subsequent to saturation and mismatch control are denoted herein as $F(u,v)$.

- (3) Unitary separable transformation is performed (either on the host or the accelerator, as negotiated):

$$f(x,y) = \frac{1}{\sqrt{H_T}} \sum_{v=0}^{H_T-1} C(v) \cos \left[\frac{(2y+1)v\pi}{2H_T} \right] \left\{ \frac{1}{\sqrt{W_T}} \sum_{u=0}^{W_T-1} C(u) \cos \left[\frac{(2x+1)u\pi}{2W_T} \right] F(u,v) \right\}$$

where: $C(u) = 1$ for $u=0$, otherwise the square root of 2 ($\sqrt{2}$);

$C(v) = 1$ for $v=0$, otherwise $\sqrt{2}$;

x and y are the horizontal and vertical spatial coordinates in the pixel domain; and

W_T and H_T are the width and height of the transform block.

- (4) Adding the spatial-domain residual information to the prediction for non-intra macroblocks to perform picture reconstruction (on the accelerator 174).
- (5) Clipping the picture reconstruction to a range of $[0, 2^{\text{BPP}}-1]$ to store as the final resulting picture sample values (on the accelerator 174).

Host v. Accelerator IDCT

As alluded to above, API 104 provides for off-host (e.g., accelerator-based) and host-based IDCT processing of multimedia content (described more fully below with Fig. 7). The transfer of macroblock IDCT coefficient data for off-host IDCT processing consists of a buffer of index and value information. According to one implementation, index information is sent as 16-bit words (although, only 6-bit quantities are really necessary for 8x8 transform blocks), and transform coefficient value information is sent as signed 16-bit words (although only 12-bits

are really needed). According to one implementation, the transform coefficient is sent as a Tcoeff data structure as follows:

```
Tcoeff {  
    TCoeffIDX (specifies the index of the coefficient in the block)  
    TCoeffEOB (denotes last coefficient in block)  
    TcoeffValue (the value of the coefficient in the block)  
}
```

TCoeffIDX: specifies the index of the coefficient in the block, as determined from ConfigHostInverseScan. There are two basic ways that TCoeffIDX can be used:

- Run-length ordering: When ConfigHostInverseScan is 0, MBscanMethod indicates a zig-zag, alternate-vertical, or alternate-horizontal inverse scan. In this case, TCoeffIDX contains the number of zero-valued coefficients which precede the current coefficient in the specified scan order, subsequent to the last transmitted coefficient for the block (or the DC coefficient if no preceding).
- Arbitrary ordering: When ConfigHostInverseScan is 1, MBscanMethod indicates arbitrary ordering. In this case, TCoeffIDX simply contains the raster index of the coefficient within the block (i.e., $\text{TCoeffIDX} = u + v \cdot W_T$)
- TCoeffIDX shall never be greater than or equal to $W_T \cdot H_T$.

TCoeffEOB: Indicates whether the current coefficient is the last one associated with the current block of coefficients. A value of 1 indicates that the current coefficient is the last one for the block, and a value of 0 indicates that it is not.

TcoeffValue: The value of the coefficient in the block. TcoeffValue shall

1 prior to any output resampling). The buffer passed to the accelerator shall contain
2 read-back commands containing a single parameter per macroblock read:

3 **MBaddress:** Specifies the macroblock address of the current macroblock in
4 raster scan order. If BPP is 8, the data shall be returned in the form of 8-bit signed
5 values, otherwise in the form of 16-bit signed values (relative to 2^{BPP-1}).

6 The data is returned to the decoder 160 in the form of (1) a copy of the
7 read-back command buffer itself followed by padding to the next 32-byte
8 alignment boundary; and (2) the macroblock data values. The macroblock data
9 values are returned in the order sent in the read-back command buffer, in the form
10 $W_T \cdot H_T$ samples per block for each block in each macroblock. Residual difference
11 blocks within a macroblock shall be returned in raster-scan order for Y, followed
12 by all 4:2:0 blocks of Cb in raster scan order, followed by 4:2:0 blocks of Cr,
13 followed by 4:2:2 blocks of Cb, and so on.

14 **Bitstream Data Buffer**

15 API 104 also supports a bitstream data buffer within operational data
16 structure(s) 204. As used herein, the bitstream data buffer, if used, primarily
17 contains raw bytes from a video bitstream to support off-host (e.g., accelerator
18 174) decoding including low-level bitstream parsing with variable length
19 decoding. According to one example implementation, the beginning of such a
20 buffer contains one or more of (1) the number '5' encoded in 8-bits to denote the
21 bitstream buffer, (2) the sequence number of the buffer within the picture, starting
22 with the first such buffer being buffer zero (0), (3) the total size of the buffer in
23 bytes, (4) if the sequence number is zero, the relative location within the bitstream
24 data of the first bit after the picture header data, i.e., the first bit of the group of
25

blocks (GOB) or slice, or macroblock layer data, and (5) reserved bit padding to the next 32 byte boundary.

The remaining contents of the buffer are the raw bytes of a video bitstream encoded according to a specified video coding format. The buffer with sequence number zero start with the first byte of the data for the picture and the bytes thereafter follow in bitstream order.

DVD Subpicture Control Buffer

As introduced above, operational data structure(s) 204 may also include a subpicture control buffer to support digital versatile disc (or DVD) applications. API 104 is invoked in support of such an application, the content of the subpicture control buffer within the operational data structure(s) 204 includes one or more of the following:

- SubpictureBufferIndicator
- ReservedBits
- BufferSize
- BlendType
- ButtonColor
- ButtonTopLeftHorz
- ButtonTopLeftVert
- ButtonBotRightHorz
- ButtonBotRightVert
- ButtonHighlightActive
- PaletteIndicator
- PaletteData
- NewSubpictureUnitSize
- DCSQTStartAddress
- SubpictureUnitData

SubpictureBufferIndicator: The number “6”, indicating a DVD subpicture buffer.

BufferSize: The total number of bytes in the buffer.

BlendType: A value of “0” indicates that no subpicture blending is active for the current picture. A value of “1” indicates that the last previously-sent subpicture data is used for blending the current picture, and a value of “2”

1 indicates that a new subpicture sent in the current buffer is used for blending the
2 current picture.

3 **ButtonColor:** Contains the color of a rectangular button on the subpicture.

4 **ButtonTopLeftHorz, ButtonTopLeftVert, ButtonBotRightHorz,**

5 **ButtonBotRightHorz:** Contains the zero-based 2-d location of the top left
6 and bottom right coordinates of the button.

7 **ButtonHighlightActive:** Indicates whether or not the button is currently
8 highlighted.

9 **PaletteIndicator:** Indicates whether or not a new palette is contained in
10 the buffer.

11 **PaletteData:** If PaletteIndicator is "1", contains the new palette.
12 Otherwise not present.

13 **NewSubpictureUnitSize:** The size of a new subpicture unit contained in
14 the buffer. If "0", indicates that no new subpicture unit is contained in the buffer.

15 **DCSQTStartAddress:** The byte location within the SubpictureUnitData at
16 which the subpicture display control sequence is found.

17 **SubpictureUnitData:** The subpicture PXD and SP_DCSQT data for the
18 new subpicture unit.

19 According to one aspect of the present invention, the control command data
20 structure and the residual difference data structure of the operational data
21 structure(s) 204 are a fixed size for each macroblock within a picture based, at
22 least in part, on one or more of the negotiated coding format, the API
23 configuration and the picture type. That is, API 104 utilizes fixed-size data
24 structures to facilitate communication between any video decoder 160 and any
25 video accelerator 174 according to any codec. Example data control command

and residual difference data structures are provided with reference to Figs. 3 and 4, respectively.

Example Data Structures

Figs. 3 and 4 graphically illustrate an example control command data structure 300 and a residual difference data structure 400 for a plurality of elements of received multimedia content. For purposes of illustration, and not limitation, the data structures are presented in accordance with the video decoding embodiment used throughout, wherein the data structures are incrementally populated with video information on a macroblock basis. According to one aspect of the present invention, introduced above, each of the control command data structures are of fixed size for each macroblock within a picture.

As shown, each element within the control command data structure 300 includes an address field 302, a pointer to an associated residual difference data structure element 304, and a command field 306. The address field 302 denotes which macroblock of a the frame the data structure element is associated with. Use of the macroblock address field 302 facilitates parallel processing of the multimedia content.

The residual difference pointer field 304 contains pointers to associated elements in the residual difference data structure 400. It is to be appreciated that not every macroblock will have residual difference data, and the amount of residual data may vary from macroblock to macroblock. Thus, use of the pointer field 304 relieves API 104 from having to inferentially associate each element of control command data structure 300 with an element of residual difference data structure 400.

1 The macroblock control command field 306 contains one or more
2 commands instructing the decoder on what action to take with respect to the
3 particular macroblock. In general, the control command field 306 contains
4 information regarding encryption of the data sent between decoder 160 and
5 accelerator 174, picture-level parameters, processing and communication control
6 parameters.

7 In addition, as introduced above, decoder 160 may well provide accelerator
8 174 with raw bitstream data, e.g., on a per-slice basis. In such an instance, API
9 104 generates a bitstream buffer to pass the raw bitstream data to the accelerator.
10 According to one implementation, analogous to the control command data
11 structure/residual difference data structure combination, the raw bitstream data
12 buffer is associated with a slice control data structure, to pass slice control
13 information from the decoder to the accelerator.

14 15 **Example Operation and Implementations**

16 As introduced above, API 104 is an enabling technology in that it facilitates
17 communication between a decoder application 160 and a hardware accelerator 174
18 as to the specific decoder/accelerator combination to be used. Having introduced
19 the architectural detail of API 104, above, attention is now directed to Figs. 5-8
20 wherein an example implementation is described.

21 Fig. 5 is a flow chart of an example method for interfacing a decoder
22 application with a hardware accelerator to cooperatively decode encoded
23 multimedia content, in accordance with the teachings of the present invention. For
24 ease of explanation, and not limitation, the method of Fig. 5 will be developed
25 with continued reference to Figs. 1-4.

Turning to Fig. 5, the method begins with block 502 which represents a step of iteratively issuing configuration commands reflecting various alternative degrees and methods of decoding acceleration capability until choosing one that is acceptable to both the decoder and the accelerator. Specifically, a media processing system element issues a ConfigInfo data structure to other media processing system elements, as the auto-negotiation process of API 104 is selectively invoked. According to one example embodiment, the auto-negotiation data structure(s) 202 of API 104 are generated by decoder 160 and reflect a proposed decoding format (ConnectMode), intermediate data format and other decoding details (ConnectConfig).

In block 504, the issuing media processing element (e.g., decoder 160) receives a response to the issued auto-negotiation data structure(s) 202 denoting whether the media processing element(s) (e.g., accelerator 174) supports the proposed media processing format defined in the auto-negotiation data structure(s) 202. If, in block 504, the proposed media processing format is not supported by one or more of the media processing elements (e.g., accelerator(s) 174), the issuing media processing element generates a new auto-negotiation data structure(s) 202 reflecting an alternate media processing configuration, block 506. In particular, decoder 160 moves to another supported media processing format and generates a ConnectMode and ConnectConfig commands reflecting the proposed media processing format. According to one implementation, decoder 160 initiates the auto-negotiation process by proposing decoding in accordance with the MPEG-2 format.

If, in block 504, the media processing format is accepted, API 104 dynamically selects one or more operational data structure(s) 204 appropriate to

1 facilitate media processing among and between the media processing elements in
2 accordance with the negotiated format, block 508. In particular, API 104 selects
3 picture parameters and buffer structures appropriate to facilitate the particular
4 media processing format agreed upon by the media processing elements (e.g., the
5 decoder 160 and accelerator 174).

6 In block 510, API 104 facilitates multimedia decoding among and between
7 the media processing elements utilizing the dynamically selected operational data
8 structure(s) 204 until the media processing has been completed. Thus, API 104
9 identifies the media processing capability of the various media processing system
10 elements, and facilitates decoding among and between these elements without *a*
11 *priori* knowledge of the particular elements used. In this regard, API 104 is a
12 ubiquitous multimedia API insofar as it facilitates communication between any
13 decoder application and any multimedia accelerator.

14 Fig. 6 is a flow chart of an example method of decoding media content,
15 according to one example implementation of the present invention. In accordance
16 with the illustrated example implementation of Fig. 6, the method begins once the
17 decoding format has been negotiated between the media processing system
18 elements, e.g., decoder(s) 160, accelerator(s) 174, etc. (block 504). The decoding
19 process of Fig. 6 begins with block 602 by saturating each reconstructed
20 coefficient value in the transform coefficient block to an allowable range. As
21 introduced above, this is commonly performed by the decoder application 160.
22 Once the saturation is complete, the saturated values are added to the coefficients
23 in the macroblock to perform mismatch control, as necessary, block 604. As
24 alluded to above, mismatch control may be necessary in MPEG-2 decoding.

1 In block 606, unitary separable transformation is performed. This
2 transformation may well be performed by the decoder application 160 on the host,
3 or by the accelerator 174. According to one innovative aspect of API 104, a
4 determination is made during the auto-negotiation process as to which element
5 will perform the transformation.

6 In block 608, the spatial domain residual difference information is added to
7 the prediction for non-intra macroblocks to perform picture reconstruction. This
8 task is typically performed off-host, i.e., at the accelerator(s) 174.

9 In block 610, the accelerator 174 performs a clipping operation to clip the
10 picture reconstruction to an appropriate range to store as the final resulting picture
11 sample values.

12 Fig. 7 is a flow chart of an example method facilitating host-based inverse
13 discrete cosine transform (IDCT), according to one aspect of the present invention.
14 In accordance with the illustrated example embodiment of Fig. 7, the method
15 begins with block 702 a determination is made as to whether the IDCT process
16 will be performed on the host (e.g., by decoder 160), or on the accelerator 174. If
17 the IDCT is performed by the accelerator, a buffer structure is established in
18 operational data structure(s) 204 of API 104 to transfer macroblock IDCT
19 coefficient data to the accelerator on a per-macroblock basis in support of the
20 transform, block 704. This process is continued until all of the macroblocks have
21 been processed.

22 If the IDCT is to be performed on the host, a first determination is made
23 whether the BPP value is greater than 8 bits, block 706. If so, the spatial domain
24 data resulting from the IDCT process performed by the decoder 160 will be
25

Effect of 8-bit Spatial Differences

PicOverflowBlocks	PicResid8Subtraction	First Pass Effect	Overflow Pass Effect (Not Allowed if Intra)
0	0	added	N/A
0	1	subtracted (no intra)	N/A
1	0	added	added
1	1	added	subtracted

1. When IntraMacroblock = 1, no overflow blocks are present.
2. When PicOverflowBlocks = 0 and PicResid8Subtraction = 1, IntraMacroblock shall be 0.

If, in block 710, the current picture is not an intra-picture then either of the 16-bit or 8-bit communication methods may well be implemented, block 716.

Deblocking Filter Control

Turning to Fig. 8, API 104 facilitates control of a deblocking filter on an accelerator 174 by the decoder 160, according to one aspect of the present invention. In accordance with the illustrated example implementation, API 104 assesses received commands for deblocking filter control commands, block 802. If deblocking filter control commands are recognized within a command received from decoder 160, API 104 generates operational data structure(s) 204 including instructions which, when received by the accelerator 174, will affect one or more deblocking filter settings, block 804. In block 806, deblocking filter control commands if present within operational data structure(s) 204, are sent for each

luminance block in a macroblock and are sent once for each pair of chrominance blocks. According to one implementation, the commands are sent in raster scan order within the macroblock, with all blocks for luminance sent before any blocks for chrominance, then one chrominance 4:2:0 command, then one chrominance 4:2:2 command if needed, then two chrominance 4:4:4 commands if needed (the same filtering is applied to both chrominance components). According to one implementation, the filtering for each block is specified by specification of the deblocking to occur across its top edge, followed by specification of the deblocking to occur across its left edge. Deblocking is specified for chrominance only once – and the same deblocking commands are used for both the Cb and Cr components. For example, deblocking of a 16x16 macroblock which contains 4:2:0 data using 8x8 blocks is specified by sending four (4) sets of two (one top and one left) edge filtering commands for the luminance blocks, followed by one set of two edge filtering commands for the chrominance. In response, to receiving such a data structure, accelerator 174 modifies zero or more deblocking filter attributes, in accordance with the received deblocking filter commands, block 808. An example data structure to effect deblocking filter commands within operational data structure 204 is provided as:

```
deblocking_edge_control {  
    DXVA_filterOn  
    STRENGTH  
}
```

DXVA_filterOn: This flag shall be '1' if the edge is to be filtered;

1 STRENGTH: This parameter specifies the strength of the filtering to be
2 performed. According to one implementation, the strength values are adopted from
3 H.263 Annex J.

4 **Alternate Implementations**

5 Fig. 9 illustrates a block diagram of a media application program interface
6 (API) according to an alternate embodiment of the present invention. According
7 to the illustrated example embodiment of Fig. 9, in addition to auto-negotiation
8 data structure(s) 202 and operational data structure(s) 204, API 900 includes
9 control logic 902, memory resources 904 and input/output (I/O) interface facilities
10 906, each coupled as shown. According to this alternate embodiment, control
11 logic 902 dynamically generate auto-negotiation data structure(s) 202, which are
12 sent to one or more media processing elements via I/O interface 906 to negotiate
13 the media processing capability of one or more media processing elements of a
14 media processing system. According to one implementation, a number of media
15 processing formats are retained in memory 904 for use in generating the auto-
16 negotiation data structure(s) 202. In one implementation, control logic 902
17 accesses communicatively coupled resources for media processing formats with
18 which to generate auto-negotiation data structure(s) 202. Control logic 902
19 iteratively issues auto-negotiation data structure(s) 202 until the elements of the
20 media processing system have agreed upon a media processing format and
21 division of media processing responsibility.

22 Once a processing format has been agreed upon, control logic 902 selects
23 one or more operational data structure(s) 204 to facilitate further media processing
24 among and between media processing elements, in accordance with the agreed
25 upon format.

Turning next to Fig. 10, a block diagram of a storage medium having stored thereon a plurality of instructions including instructions which, when executed, implement the teachings of the present invention, according to yet another embodiment of the present invention. In general, Fig. 10 illustrates a storage medium/device 1000 having stored thereon a plurality of executable instructions 1002 including at least a subset of which that, when executed, implement the adaptive API 104 of the present invention. When executed by a processor (132) of a host system (100), the executable instructions implementing API 104 identify and characterize the processing capability of a multimedia processing system, and dynamically adjusts one or more operational settings to operatively interface any decoder application with any multimedia accelerator. In this regard, API 104 is an extensible, universal multimedia API. According to one implementation, API 104 selectively modifies one or more operational settings to improve multimedia processing performance of the host system (100) based, at least in part, on the identified functional capability of the one or more elements of the multimedia processing system.

As used herein, storage medium 1000 is intended to represent any of a number of storage devices and/or storage media known to those skilled in the art such as, for example, volatile memory devices, non-volatile memory devices, magnetic storage media, optical storage media, and the like. Similarly, the executable instructions are in machine language, interpreted languages, and/or other source code that will be interpreted, such as, for example, C, C++, Visual Basic, Java, Smalltalk, Lisp, eXtensible Markup Language (XML), and the like. Moreover, it is to be appreciated that the storage medium/device 1000 need not be co-located with any host system. That is, storage medium/device 1000 may well

1 reside within a remote server communicatively coupled to and accessible by an
2 executing system. Accordingly, the software implementation of Fig. 10 is to be
3 regarded as illustrative, as alternate storage media and software embodiments are
4 anticipated within the spirit and scope of the present invention.

5 Although the invention has been described in language specific to
6 structural features and/or methodological acts, it is to be understood that the
7 invention defined in the appended claims is not necessarily limited to the specific
8 features or steps described. Rather, the specific features and steps are disclosed as
9 example forms of implementing the claimed invention.

10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25